CLAIMS

We claim:

1. A method by which a second system maintains connections for a failed first system, comprising:
receiving ownership information from the first system on which an application is running;
determining that the first system is in a failed state; and
assuming a connection for the first system.

2. The method of claim 1, wherein the ownership information is received when the second system is booted.

3. The method of claim 1, wherein the ownership information comprises:
a fail-over policy;
a set of one or more IP addresses owned by the system;
a range of port numbers;
an application that is currently running on the system;
a ARP ownership policy;
a current protocol;
a set of one or more MAC addresses tied to the IP address(es); and
a cluster node ID.

4. The method of claim 1, wherein the first system and the second system solicit ownership information from each other.

5. The method of claim 1, further comprising
determining that the first system has returned from the failed state to a normal state; and
returning responsibility for new connections to the first system.

6. The method of claim 1, wherein the first system and the second system are within a cluster of systems.

7. The method of claim 1, wherein the connections are TCP connections.

8. The method of claim 1, further comprising snooping connection data of the first system.

9. The method of claim 1, further comprising continuing the application from the point at which the first system failed.

10. The method of claim 1, further comprising:
attempting by a third system to assume the connection; and

sending from the third system to the second system a request to assume the connection.

11. A system for maintaining a connection within a network, comprising:
means for broadcasting ownership information between a first system on which an application is running to at least a second system within the network;
means for determining that the second system will assume the connection for the first system if the first system fails;
means for transmitting to the second system, packets sent to and received by the first system; and
means for determining that the first system is in a failed state;
means for continuing the application on the second system from the point at which the first system failed.

12. The system of claim 11, further comprising a means for broadcasting ownership information when each system is booted.

13. The system of claim 11, wherein the ownership information comprises:
a fail-over policy;
a set of IP addresses owned by each system within the network;
a range of port numbers;
an application that is currently running on each system;
a ARP ownership policy;
a current protocol;
a set of MAC addresses tied to the IP address(es); and
a cluster node ID.

14. The system of claim 11, further comprising:
means for returning the first system from a failed state to a normal state; and
means by which connection is regained by the first system from the second system.

15. An article for maintaining connections by a second system for a first system, comprising:
a computer-readable signal-bearing medium;
means in the medium for receiving ownership information of the first system;
means in the medium for determining that the first system is in a failed state; and
means in the medium for assuming a connection for the first system.

16. The article of claim 15, further including:
means in the medium for continuing an application on the second system from the point at which the first system failed.

17. The article of claim 15, wherein the ownership information of the first system is broadcast when the system is booted.

18. The article of claim 15, wherein the ownership information comprises:
    a fail-over policy;
    a set of IP addresses owned by the first system;
    a range of port numbers;
    an application that is currently running on the first system;
    a ARP ownership policy;
    a current protocol;
    a set of MAC addresses tied to the IP address(es); and
    a cluster node ID.

19. The article of claim 15, further comprising:
    means in the medium for determining that the first system has returned from a failed state to a normal state; and
    means in the medium for returning responsibility for connections to the first system.